

THE USE OF PACKETIZED TELEMETRY IN INVERSE T1 MULTIPLEXING

**Jason Urban
Apogee Labs, Inc.
North Wales, PA**

ABSTRACT

As the number of telemetry applications at sites increases, the need for a higher bandwidth link from site-to-site grows. The use of an inverse T1 multiplexer allows the site to utilize multiple T1 lines rather than more costly higher bandwidth lines. There are many advantages to using a recognized packet standard, such as IRIG-107/98, over simply routing the streams through T1 lines. One advantage is that there is a total separation of data stream clock from T1 transmission clock, reducing synchronization circuitry and overhead. Another advantage is that the use of packets also allows for a smooth reconstructed clock phase on the receiving site, creating a virtually seamless transmission of clock and data. And, finally, by using a recognized packet standard, the inverse T1 multiplexer may easily be integrated into ranging and telemetry systems already incorporating packetized telemetry. This paper will discuss the combination of packets and inverse multiplexing to achieve an expandable transmission system capable of delivering a high bandwidth data stream across multiple T1 lines.

KEYWORDS

Inverse Multiplexing, Packetized Telemetry, IRIG 107/98, CCSDS, T1, DS3

INTRODUCTION

Telemetry sites have an ever-increasing need for high bandwidth data transmission. They need to move data, such as digital telemetry, analog, and video, from one station to another. Along with increased data rates comes increased transmission line cost. As the data rate approaches 1MHz, T1 lines are a logical choice. But when the data rate exceeds a T1 rate, the next common publicly available connection is a DS3 line. Although the 45Mbps DS3 line can handle most simple telemetry applications, the cost and maintenance of a DS3 line can quickly exceed budgets. Inverse T1 multiplexers exist to alleviate the cost issue of moving data across a DS3 line[5]. Packetized telemetry, such as the IRIG 107/98 standard, offers a method of data transmission already in use by the telemetry world[4]. The combination of packetized telemetry and inverse multiplexing

provides a data transmission method with all elements familiar to those involved with telemetry.

INVERSE MULTIPLEXING

Many sites implementing data communications are limited to a single data line. If a site has a 64kps line, then that site may only transmit or receive data at the 64kps rate. This does not take into account any overhead for the transmission protocol, thereby reducing the actual data rate to less than 64kps. As the site's data rate increases, the need for faster communication lines becomes an issue. However, not only does the data rate increase, so does the cost of line leasing, maintenance, and equipment. This cost increase grows exponentially. The user and site are then restricted not only by the bandwidth of the current line and but by the cost of upgrading to a faster line.

Inverse multiplexing is the solution to the need for higher bandwidth at a low cost. Standard multiplexing techniques combine multiple channels into a single, higher rate composite stream. Inverse multiplexing, in effect, does the opposite of standard multiplexing. Rather than a single composite stream, a single stream is split into multiple, lower rate streams. These low rate streams may then be transmitted over multiple lines. At the opposite end, the inverse demultiplexer combines the multiple streams. The result of this combination is a single stream at a much higher bandwidth. Figure 1 is an example of an inverse multiplexer to transmit a high rate PCM stream over multiple T1 lines.




Figure 1 – Inverse T1 Multiplexer Example

As the data rate to an inverse multiplexer changes, the number of T1 lines needed to complete a system can change. T1 lines are added to the system to facilitate a rising input rate and are removed as the input rate decreases. This addition and subtraction of T1 connections is known as dynamic bandwidth allocation. Dynamic bandwidth allocation allows for a changing input rate without disturbing the original T1 connections. By measuring the input rate, the system calculates the number of T1 channels needed to carry the data. The number of connections is automatically adjusted if the system knows the maximum number of T1 channels. By only using the minimum number of connections required for the input rate, costs may be reduced by not having a large number of connections for a low input rate.

From end to end, T1 lines can take many paths. One line may travel from a ground station through a satellite back to a ground station, while another line may travel on the ground but through several switched networks. In the case of these two lines, the satellite delay is relatively long compared to the ground based line. If these two connections and a third point to point connection is used in an inverse multiplexer system, the data carried by the point to point connection would arrive before the other two connections as shown in Figure 2. System complexity increases because these delays must be taken into account when reordering the data.

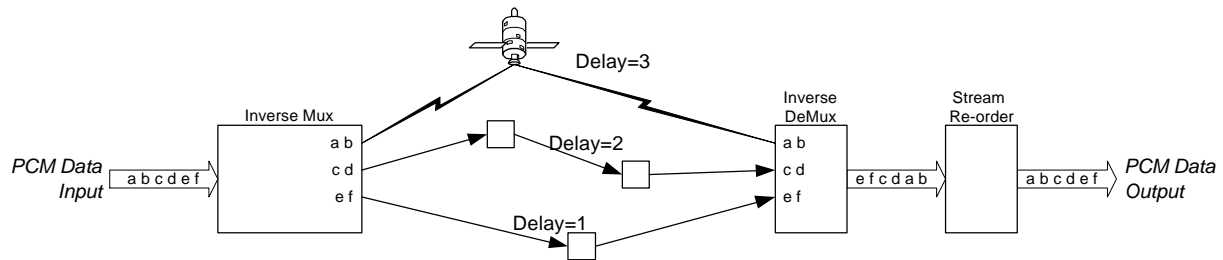


Figure 2 – T1 Line Delays

CCSDS AND IRIG 107/98

Packetized telemetry may be split into several basic layers, as shown in Figure 3. This figure represents both transmission and receiving ends of a telemetry system using the CCSDS format. In this example, four sources of PCM data are received and formatted into source packets. Source packets are of variable length and are an even number of octets. It is within these source packets that the application ID, for example, is defined. If time-tagging is required, it is also included in the source packet. In order to group the data by application, such as video and tracking data, source packets are grouped into virtual channels. All virtual channels are multiplexed into a master channel to prepare the data for transmission over a single physical link. Before passing to the physical channel, the master channel data is formatted into a transfer frame. Transfer frames are fixed in length throughout the mission and carry the channel data as well as error correction and channel identifiers[1].

In order to simplify the CCSDS approach for use in an already complex inverse multiplexing system, a similar approach to what is used in standard multiplexers may be taken. When using a low number of similar rate channels, multiple virtual channels are no longer needed. A single virtual channel combined with a single master channel effectively reduces the amount of hardware needed to implement the packet telemetry system[2]. As shown in Figure 4, the packetizer layer gathers source packets by means of a data collector bus and formats the data into a transfer frame. The end user of the data simply receives the stream, depacketizes, and sends the source packets to the appropriate output channel[4].

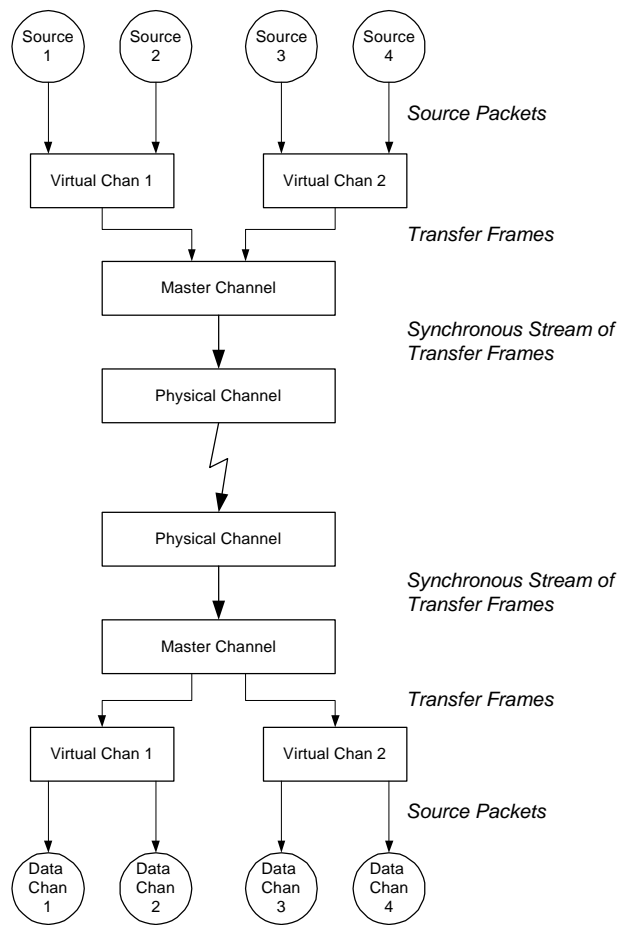


Figure 3 – IRIG 107/98 Data Layers




Figure 4 – Simplified Packet Example

A similar standard to CCSDS that is now being used on the telemetry ranges is IRIG 107/98. The IRIG 107/98 standard was originally intended to provide an on-board flight record standard. When used in the application of inverse T1 multiplexing, IRIG 107/98 makes no changes to the CCSDS standard of defining source packets and transfer frames. The IRIG 107/98 standard is used in this application to show that a record standard can be used in point-to-point communications[3].

USE OF IRIG 107/98 IN AN INVERSE T1 MULTIPLEXER

When the need for higher bandwidth data arises, common T1 lines may be used. Many sites are connected to one or more T1 lines, each line with a 1.544Mbps rate. Rather than stepping up to a more costly DS3 line, with a 45Mbps rate, an inverse T1 multiplexer may be used to spread the higher rate channel over multiple T1 lines.

A T1 frame, also referred to as a DS1 frame, was originally conceived to carry 24 voice channels. Today, each of these 24 channels is commonly used to transfer data and channels may even be combined to transfer one channel of data. The T1 frame contains 24 eight bit channels and one frame bit(F-Bit) for a total of 193 bits per frame. T1 transmits 8000 frames per second for a composite rate of 1.544Mbps. Due to the F-Bit, there is 8kps of overhead in T1 transmission. In this application of an inverse T1 multiplexer, we will combine all 24 channels in the T1 frame for data transmission[5].

Primary Header	Secondary Header	Data Field
ID, Seq#, Length	PCM rate, #active chan	PCM Data

Figure 5 – IRIG 107/98 Source Packet

By using packetized telemetry techniques, data may be carried across T1 lines in a known format. The full IRIG 107/98 source packet for this application is shown in Figure 5. This frame is 1916 bytes long, with the primary header 6 bytes, the secondary header 4 bytes, and the data field being 1906 bytes. Using the standard set by IRIG 107/98, the primary source packet header is shown in Figure 6 and the secondary header is shown in Figure 7. For details about fields contained in both the primary and secondary sources, see reference [1]. This paper will only discuss fields relevant to the use of IRIG 107/98 in an inverse T1 multiplexer.

Version Number	Packet Identification			Sequence Control		Data Length n-1 bytes
	Type Indicator	Secondary Header Flag	Application ID	Grouping Flags	Sequence Count	
3 bits	1 bit	1 bit	11 bits	2 bits	14 bits	16 bits
000	0	1	0011001cccc	00	counter	1905d

Figure 6 – Primary Header

Within the primary header is a secondary header flag. By setting this flag to a one, the depacketizer knows that a secondary header is present. The lower four bits of the Application Process ID field contain the T1 channel number. Each T1 channel connected to the inverse T1 multiplexer is numbered

PCM Frequency	# of Active Channels	Spare
25 bits	4 bits	3 bits

Figure 7 – Secondary Header

starting from one. This channel number, in conjunction with the sequence number, allow the data to be properly reordered at the T1 receive end.

Because the original PCM rate is embedded in the secondary header, a near-original clock rate may be constructed. The use of a Direct Digital Synthesizer(DDS) allows for fine tuning of the output clock. A baseline frequency is established by reading the frequency from the secondary header. Slight adjustments are determined by buffer positions to prevent under- or over-flow during output. Because a DDS does not change phase when adjust frequency, a smooth clock output phase is achieved.

In order to keep the system keyed off a particular event, a Sample Interval signal is established. The Sample Interval is used to key off events when generating source packets. During each Sample Interval, a number of source packets are sent to the T1 Line Interface memory.

A SYSTEM DESIGN EXAMPLE

In order to better explain the use of packetized telemetry in an inverse T1 multiplexer, this paper will explain a typical system.

By setting up a Sample Interval to occur every 10ms, a system heartbeat is established. This heartbeat keys the process of memory transfers and packetizing. During a Sample

Interval, the input stream is sent through a serial to parallel converter and the data stored in a ping-pong buffer. The other side of the ping-pong buffer contains a count of data words per Sample Interval along with the actual data. The packetizer uses this count to determine the number of T1 channels needed and to determine the packet size for partial packets when there are not enough data words to fill the maximum data packet size of 1906 bytes.

Independent of the Sample Interval, the T1 Line Interface retrieves source packets from the T1 Line Interface memory. Each source packet is inserted into the T1 frame as unchannelized data. If a source packet does not fill a T1 frame, the T1 framer adds fill data to the end of the source packet. The frame bit is added to the stream and the new frame is sent out to the physical T1 line. At the T1 receiving side of the system, each T1 frame is received and the source packet extracted. This source packet is stored in one side of a ping-pong buffer for the depacketizer to read.

Due to T1 channel delays, the depacketizer must first determine which line has the longest delay. Once all connections have been established and T1 frames are being received, the depacketizer polls all channels for the lowest sequence number. All channels are then read by the depacketizer in sequential order starting with channel one and the packet containing the sequence number determined by the line with the longest delay.

Some fault tolerances must be established at this point in the system. Many source packets will be lost resulting in data loss if a T1 line were to be dropped during a data transmission. In this case, the depacketizer must determine that a channel is missing. By adding fill data in place of the missing T1 channel, there will be data errors, but the sequence of source packet reading will remain. Data from other channels will continue to be depacketized because the sequencing will remain intact. An occasional source packet loss or source packet corruption will result in bit errors but not necessarily a complete data loss as in the last case. The depacketizer and associated processor determine errors in the source packet length that would be the result of T1 bit errors. By forcing this length error to a fixed value, the depacketizer only introduces bit errors rather than sequence errors or complete data loss.

As each source packet is depacketized, the data is stored in an output buffer. This output buffer is setup as a FIFO, with flags at empty, 50% full, and 100% full. In order to reconstruct the output clock, the output system needs to determine the original data rate. Each T1 channel contains source packets which contain the original clock frequency. The output system then has multiple sample of the original rate. By a majority voting system, the output system can setup a DDS for the correct output rate.

Using the 50% full flag from the output buffer, the output system can vary the output frequency to keep the output buffer from under- or over-flowing. Whenever the 50% full flag is set, the parallel to serial converter will output data until the buffer drops below 50% at which time the clock rate will decrease slightly to prevent gaps. The buffer will then approach the 50% mark and the clock rate will return to the original rate. This variance in clock rate will allow for a smooth output clock with only minor rate variations.

CONCLUSION

It has been shown that the use of packetized telemetry in an inverse multiplexer has several advantages. By utilizing the IRIG 107/98 standard and common T1 lines, high bandwidth data transmission using DS3 lines can be avoided. The use of the secondary header and a DDS allows for a smooth clock and data reconstruction. Data sequence is maintained by the packetizer using the sequence number in the primary header for delay consideration and error free data transmission. The IRIG 107/98 and CCSDS standards are currently in use and allow for an effortless integration of a packetized inverse T1 multiplexer into telemetry systems. These standards also allow the decoupling of input bit rate from T1 channel rate.

REFERENCES

- [1] Packet Telemetry, CCSDS 102.0-B-4 Blue Book, November 1995
- [2] Packet Telemetry Services, CCSDS 103.0-B-1 Blue Book, May 1996
- [3] Telemetry Group Range Commanders Council IRIG Standard 107/98, June 1998
- [4] International Telemetry Conference Volume XXXIII, 97-02-3, 'Multiplexer/Demultiplexer Implementation Using a CCSDS Format', David L. Grebe, Apogee Labs, Inc., 1997
- [5] Spohn, David L., Data Network Design, Second Edition, McGraw-Hill, 1997